SF 298 MASTER COPY KEEP THIS COPY FOR REPRODUCTION PURPOSES

REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	9/30/98	Final Report, 6/15/95-6/30/98

4. TITLE AND SUBTITLE

Algorithms for Image Compression, Distributed Communication Networks and Distributed Resource Allocation

5. FUNDING NUMBERS

DAAH04-95-1-0101

6. AUTHOR(S)

Basilis Gidas
Sumit Ghosh

7. PERFORMING ORGANIZATION NAMES(S) AND ADDRESS(ES)

Brown University
Providence, RI 02912

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

U.S. Army Research Office
P.O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSORING / MONITORING AGENCY REPORT NUMBER

ARO 34383.5-MA-FR I

11. SUPPLEMENTARY NOTES

The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.

12a. DISTRIBUTION / AVAILABILITY STATEMENT

Approved for public release; distribution unlimited.

12 b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

During the period of the grant, 6/15/95-6/30/98, we investigated problems in two major areas: (i) Image Analysis Tasks, including image interpretation, object recognition, and tracking; (ii) Distributed communication networks and distributed resource allocation. In the former area, we developed and explored three paradigms: hierarchical compositional models; deformable templates for a host of applications including medical tasks; and HMM/deformable templates for tracking and recognition of moving objects. In the later area, we developed a novel scheme for managing buffer overflows; a new framework for network security; and a mathematical framework for synthesizing distributed algorithms.

14. SUBJECT TERMS

Object recognition, tracking, context-free-grammars, templates, buffer management, ATM networks, distrubted algorithms.

15. NUMBER IF PAGES

11

16. PRICE CODE

17. SECURITY CLASSIFICATION OR REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

TITLE

Algorithms for Image Compression, Distributed Communication Networks and Distributed Resources Allocation

FINAL REPORT

AUTHORS

Basilis Gidas (Professor)
Sumit Ghosh (Professor)

DATE

June 15, 1995 – June 30, 1998

U.S. ARMY RESEARCH OFFICE

GRANT NUMBER

DAAH04-95-1-0101

INSTITUTION

Brown University

TABLE OF CONTENTS

FINAL REPOTRT

Summary of Research

Our research supported by the ARO/FRI grant has had several successes in its two principal lines of investigation: (i) Algorithms for Image Analysis Tasks such as image interpretation, object recognition, tracking, and image coding; (ii) Networking, Distributed Resources Allocation. and distributed Algorithms. The algorithms have been strongly motivated by, and have been implemented on, real-world applications. Below summarize some of major contributions.

1 Algorithms for Image Analysis

Compositional Models for Object Recognition: Objects in an image present themselves in a variety of poses, illumination, and states of contiguity and occlusion. The task of an artificial vision system is to segment *and* recognize, and the challenge is in achieving an adequate articulation of these two operations, the former being classically regarded as a low-vision problem while the latter as a high-vision problem. In most real-world applications, it is impossible to perform segmentation in the absence of high level models, just as it is impossible (in the presence of high levels of local ambiguity) to recognize objects without segmenting them from the background and from each other. As result invariant object recognition should be addressed within the context of hierarchically structured models. We introduced a class of such models we call *compositional models.* They are based on the observation that objects can be decomposed into parts which, in turn, can be further decomposed recursively into still smaller components, thereby obtaining a hierarchy of entities, e.g. discontinuities, contours, strokes, surfaces, cylinders, and , at the highest levels, objects such as pieces of computers or airplanes. This recursive decomposition, and the understanding of the relation between the parts and the whole and how this changes when factors that determine object appearances vary, is the basis of the compositional representation.

Object recognition is in some sense the inverse of this process; starting from the gray-level image data, local entities are recursively composed, subject to syntactic contextual constraints articulated by *compositional rules*, to form higher level constituents, giving rise to *labeled trees.* Then a probability distribution on the ensemble of all labeled trees is devised using ideas from coding theory and the MDL (Minimum-Description-Length) principle. The resulting probabilistic system includes probabilistic Context-Free and Context-Sensitive Grammars. A global tree typically has a large number of nodes labeled by objects or object components at various levels of the hierarchy; while the composition rules can, in principle, access all the information encoded at lower- levels of the representation, computations are made practical by the fact that at in any given application only a judiciously chosen subset of low- level nodes communicate with higher-level parts of the representation.

Compositional type procedures were developed for the following three problems:

(i) A Context-Free-Grammar (CFG) type compositional procedure was developed in for detecting multiple 2-D convex objects in images where background noise conspires to make regions of high structure which need to be distinguished from more regular shapes that represend the real objects. The procedure was succesfully applied to the detection of mines in shallow water optical imagery. The basic ingredients of the approach are as follows:

A data model was devised to assign a likelihood to the grey-level data collected in the vicinity of an object. More specifically, given any boundary outlining a simly connected region, the distribution of the data within the boundary and within an annulus of the surrounding of the boundary was modeled in such a way as to favor contrast between the intra-boundary and extra-boundary pixels. Additionally, a "background" or "null" model was built on the assumption

of independence and homogeinuity of the pixel grey-levels regardless of the placement of the perported boundary.

The goal of the object dedection algorithm was to perform a proper hypothesis test at each location in the image, testing the hypothesis "location is part of the object" versus the hypothesis "background".Although the test procedure, to be described shortly, could be executed quite efficiently, the algorithm was made much faster by "preprocessing" the image in order to flag potential positions of interest. The procedure was designed to produce multiple "hits" for each genuine object, so as to insure that no object was missed.

The "object present" hypothesis is highly compound, involving essentially all convex shapes. Shape, then, amounts to a nuisance parameter, which calls for an application of the generalized likelihood ratio test. This involves the daunting task of evaluating the data likelihood at the *maximum estimator of the boundary.* Determining the most likely placement of the boundary, under the "object present" hypothesis, is thus the essential issue of the detection algorithm.

For this purpose, we devised CFG type compositional rules which generate convex shapes only: Consider two neighboring sectors with equal angles amanating from a location of interest (a "flagged" pixel). Associate with each sector a portion of the boundary of a convex region that contains the flegged pixel. Allow theses sectors to join into a larger sector provided that the combined boundary is itself a portion of a boundary of a convex region containing the flagged pixel. The syntactic constraint is easy to work out, and the recursive appliaction of the rule eventually creates a closed contour which is necessarily a convex region. In fact all convex regions can be generated in this manner.

This representation of (digital) convex regions is equivalent to a CFG, and can be exploited to *exactly* solve the problem of finding the most likely convex boundary of a purpoted object. This turns out to be a dynamic programming problem which was solved by a multilevel (coarse-to-fine) procedure. The maximum likelihood boundary replaces the nuisance parameter and the likelihood ratio is measured. A threshold determines the false/missed detection tradeoff.

(ii) A prototype upper-case handwritten character recognition system was developed based upon a compositional hierarchy with a prior distribution. The characters were drawn on a pad with a stylulus whose position was sampled at constant rate. The data is simply the collection of sampled locations (information about the order of locations was ignored). A hierarchical representation was designed for the objects in an object library; the library includes includes upper-case letters, but also numerous other "objects" that emerge from intermidiate-level representations, including "lines", "arcs", "T-junctions", etc. The hierarchy was generated via syntactic rules under which entities are composed to form composite entities. For example, two pixels are composed into a mini-stroke ("linelet") if their distances are below a certain threshol. Approximately fifty rules were used.

The probability disttribution was devised by exploring the MDL principle. The idea was to encode an object hierarchy as if it were to be transmitted over a channel or stored on a disk. A "sensible" encoding would assign shorter codes to intuitively succinct descriptions, such as nearly collinear points as a line segment versus their description as individual and independent locations. An optimal interpretation is an assignement of objects that minimize the description length; a greedy algorithm was devised for obtaining an approximate solution of this minimization problem. The performance of the procedure is competative to, if not better than, that of commercial systems.

(iii) A hierarchical/syntactic system was developed for recognizing objects in a small database of 2-D industrial tools, in images highly degrated by noise, blur, occlusion, and overlap. The system has three basic components: (a) An object representation with two levels of hierarchies. The top level views objects as a concatenation of of its articulated joints and parts; it is repre-

sented by a -graph. Each node in the graph is associated with a "high-level" object that may be a common component of several objects; the arcs correspond to synctatic constraints relating the high-level componets. This top level can be expressed as a formal CFG. The lower level of hierarchy serves to represend top level objects as a concatenation of lower-level primitives. The comcatenation process is described by a Markov process inspired by the conceptual ingridients of the *Polygonal Markov random Fields.* (b) The lower-level elementary units intract with local descriptors of the grey-level data; these descriptors are designed in terms of nonparametric statistic (ranks and Smirnov-Kolmogorov statistics). (c) The combination of (a) and (b) leads to a formulation of the recognition problem as a global optimization problem which, in view of the recursive structures, lends itself to variations of Dynamic Programming (DP). The DP involves a large state space, and it requires the maintance of a multitude of intermidiate data structures. This required the development of efficient search strategies and pruning techniques.

Deformable Templates and Digital Anatomy: The deformable template methodology represents objects or other complex global structures in terms of prototypes (templates) and rigid and/or elastic-like deformations of the prototypes. Computational engines have been developed for identifying an optimal transformation that brings a given template into a target image; the resulting *image labeling* can be assesed via the *image likelihood*, giving a statistically-sound foundation for image analysis, understanding, and interpretation. The computational algorithms have explored "jump-diffusion" processes, and their implementations are hierarchical, multi-staged (multi-grid), or use appropriate *wavelet* expansions.

The methodology of deformable templates has been applied to a host of applications [?, ?, ?] including: (i) The detection and recognition of subcellular shapes (such as membrances and mitochondria) in electron migrographs, (ii) the representation of the anatimical variation of the geometry and "shape" of 2-D surfaces in the brain– specifically, of the cortical and hippocampal surfaces in an ensemble of Macaque monkeys; (iii) the tracking and recognition of aircrafts, based on the fusion of narrow-band radar (for detection) and high resolution radar (for recogntion); and (iv) Digital anatomy. Next we briefly describe (i) and (iv).

(i) Electron micrograph images exhibit three basic kinds of variability: *structure or shape variability* associated with shape and scale variability of various subcellular structures, e.g. oranelles and amoebae; *internal constituent variability* of the textures making up the shapes; and *complexity variability* due to the varying number of shapes-the number not known a priori.

The shape variability is articulated by defining *organelle templates* corresponding to fixed graphs with *rigid regularity* forcing global connectivity of the boundary.The teplates are made flexible by attributing them with geometric group operators (scale, rotation, translatin) resulting in rabber-band like transformations of the rigid templates. The *textual variability* and sensor noise are accomodated via likelihoods costructed from Markov Random Fields models of the textured regions. The *complexity variability* is accomodated via the construction of a *space of scenes.* A scene consists of a number of template shapes together with their associated transformations. Choosing a scene corresponds to choosing both the parameters determining the shapes as well as the *model order* most representative of the data. The computational algorithms involves the solution of stochastic partial differential equations ("diffusion" equations), and have been implemented on a massively parallel SIMD architecture; they involve also a discrete part-the so called "jump"-which accounts for the variable number of shapes.

(iv) Imaging modalities such as MRI, CT, and PET, provide detailed information regarding the anatomy and physiological function of specific subjects. But the interpretation of the data has been hindered by the inability to expiditiously relate the information to specific anatomic regions–a problem central to the Visible Project undertaken by the National Library of Medicine.

We developed algorithms to transform the morphology represented in atlases to observed instances of morphology as observed from a particular imaging modality in a particular patient; thereby transmitting relevant information from atlases to a particular image in clinical setting. This work involved templates of low complexity, but designing templates for human anatomy is a task of orders of magnitude higher.

Simultaneous Tracking and Recognition: We formulated and implemented a coherent framework for the simultaneous trcking and recognition of moving objects, on the basis of dynamically changing sequences. The framework has three basic components; a deformable template representation of objects; dynamical equations of motion derived from Lagrangian mechanics; and an observation model designed by using nonparametric image processing techniques. The last two componets combined together give rise to a HMM which is equivalent to a non-linear filtering problem. This problem is solved by a new iterative algorith based on Monte Carlo sampling and statistical resampling.

The moving objects may be rigid or deformable. In the former case, shape variability involves changes in location, rotation, and scale. In the deformable case, variability involves both global and local deformations. The parameters associated with these deformations are viewd as "generalized" coordinates, and thier equations are driven by random generalized forces and torques. The framework has been succesfully implemented on an number of enviroments highly degrated by clutter, overlap, and other degradations.

2 Distributed Networks and Resource Allocation

Buffer Management in Networking: Fundamentally, the inclusion of bursty traffic, the desire to achieve a cost-effective network design, i.e. to obtain maximum throughput in a given network while minimizing cell loss, and the limitation of maximum buffer size, derived from basic digital design principles, lends significant importance to the issue of buffer management in cell switching networks including ATM networks. The buffer management literature is rich and includes reports on improved network performance through priority schemes, fixed thresholds, admission control, buffer partitioning, fuzzy thresholds, fuzzy rule-based congestion and admission control, fuzzy policing mechanisms, and a shared memory buffer for all output ports of a switch subject to a delayed pushout scheme. While highly effective, these techniques suffer from two key weaknesses. First, they are susceptible to cell loss due to buffer overflow. Second, the performance reported for all of these schemes have been obtained for a single switch and, consequently, the results are not necessarily applicable to a realistic cell-switching network consisting of multiple switches. This research presents a novel scheme – Guaranteed-No-Cells-Dropped (GNCD), that eliminates cell loss due to buffer overflow. In this approach, the switching node first records the current buffer occupancy and then determines the absolute number of empty slots. It then admits from the input cell-burst from other switches an exact number of cells equal to the number of empty slots. The remainder of the cells are blocked at the sending switch. The computation involved in the decision to admit or refuse entry into the buffer is simple and fast – a key advantage for high-speed networks. Furthermore, buffer overflow is eliminated implying no cell loss – a key potential advantage for specific traffic classes in ATM. Although the fuzzy thresholding-based buffer management scheme holds a significant potential, according to the literature, to effectively reduce cell loss due to buffer overflow, GNCD achieves in eliminating cell loss due to buffer overflow completely. Although the fate of a large fraction of the blocked cells, caused by excess traffic, is similar to that in the fuzzy scheme, it is logical to assume that the prevention of cell loss due to buffer overflow manifests through increased

blocking. GNCD includes an effort to reroute the cells, constituting the increased blocking, to their eventual destinations. Clearly, these rerouted cells may incur delays and cause other cells in the network to suffer additional delays. Thus, the absolute current buffer occupancy and the size of the incoming cell-burst constitute the basis for admitting or blocking cells into the buffer. The GNCD scheme is modeled for both (1) a single, stand-alone switch to facilitate a meaningful comparison with other competing approaches in the literature, and (2) a 50-switch, representative cell-switching network to study its performance under realistic conditions. Both models are simulated, the first on a uniprocessor computer and the second on a testbed consisting of a network of 25+ Pentium workstations under linux, configured as a loosely-coupled parallel processor. For the simulation of model 1, a representative traffic stimulus, characterized by a two-state, "on-off" Markov Chain model and exponentially distributed departures, is utilized. A comparative analysis of the results reveal that the GNCD scheme yields a slightly higher throughput than the fuzzy scheme while ensuring zero cell loss due to buffer overflow, over a range of values of input cell arrival rates. Performance results also indicate that, contrary to the general expectation that large buffers will enhance throughput and mitigate congestion, the throughput and blocking behavior of the GNCD scheme are unaffected by increasing the buffer capacity beyond a "critical buffer size." For the simulation of model 2, the input traffic distribution consists of a total of 1.0 – 1.5 million cells that are asserted into the network for each experiment, with the video traffic synthesized from an actual movie in MPEG-1, the audio traffic obtained using an ON/OFF model utilizing parameters extracted from an actual voice segment, while the generator for the data traffic utilizes a Markov chain model, subject to traffic shaping. The results corroborate those in model 1 in that, even for a large-scale representative cell-switching network, the throughput is higher and blocking is lower under GNCD than for the fuzzy thresholding approach, over a range of input traffic distributions corresponding to throughputs of 73% to 91%, although the difference in the throughputs and blocking are negligibly small for both low and high values of input cell arrival rates. At either extremes of the input cell arrival rate values, the behaviors of both GNCD and fuzzy thresholding are similar. However, while the cell loss due to buffer overflow is nil, the average delay incurred by the cells in the network are higher for GNCD than the fuzzy approach, for the same range of input traffic distributions, by a factor ranging from 1.43 for relatively light input traffic density to 1.92 for moderate input traffic to 0.93 for heavy traffic. This research notes that GNCD incurs less blocking than the fuzzy scheme and that the fuzzy scheme incurs cell loss through buffer overflow, and argues conservatively that the cells lost through buffer overflow in the fuzzy scheme manifest through increased blocking at the sending switches and that effort must be expend to reroute them to their respective destinations. Under these circumstances, the average delay incurred by the cells in the network is higher for GNCD than the fuzzy scheme by a factor ranging from 2.5 for relatively light input traffic density to 4.72 for moderate input traffic to 1.6 for heavy traffic.

Security on Demand in ATM Networks: The traditional, certification based approach attempts to guarantee the security of a network through risk analysis that involves the static analysis of the risks, threats, resource availability, cost, and other parameters. Certified networks are generally confined to a limited domain and, as a result, they are isolated, costly, and under-utilized. This paper presents a fundamental framework for network security that realizes a unified and comprehensive view of security among the civilian, military, and government networks. It also offers a unique set of principles, inherent in the framework, to enable the realization of the dynamic risk assessment approach across the different types of networks and for any type of user – civilian, military, or government. The framework approach has been

adopted by the National Security Agency as the basis for its Network Rating Model. Under dynamic risk assessment, security is viewed as a resource and the network security problem is mapped into a resource allocation problem that takes into consideration, dynamically, the risks, threats, resource availability, cost, and other parameters. Thus, this approach reflects a pragmatic, cost effective, and best effort approach. It also derives a part of its momentum from the increasing trend towards the merging of the civilian, military, and government networks. The framework consists of eight perspectives and nine attributes. The perspectives, termed pillars in this paper, individually provide orthogonal views of network security and collectively constitute a comprehensive stable structure that supports the total network security. The attributes refer to the inherent characteristics of a secure network. While the details of the attributes and the pillars are specific to the network in question and are likely to undergo evolution as the network technology matures, the framework provides a basis to address, fundamentally, every weakness in a given network. Furthermore, it applies to every level of the network, starting at the highest network of networks level and down to the single computing node that maintains connections with other nodes. Thus, the framework enables the understanding of the security posture of an individual network, in a comprehensive manner, the comparative evaluation of the security of two or more networks, and the determination of the resulting security of a composite network that is formed from connecting two or more networks with known security.

In combination with the fundamental unique characteristics of ATM networks namely, the call setup process, the framework also offers "security on demand," a new paradigm in network security. The proposed scheme is consistent with the basic characteristics of and is fundamentally the most logical approach to security in ATM networks. Unlike in data networks such as the Internet, where the actual nodes participating in routing a user's message are unknown a priori, in ATM networks, under call setup, a virtual path for a user's message is determined through negotiation prior to launching the traffic. The negotiation generally includes the user bandwidth and Quality of Service (QoS) request. The principles underlying the framework enable the inclusion of the user requested security in the dynamic, negotiation process. Thus, this unique, two-step process ensures that before any user traffic is propagated, the entire path from the source to the destination ATM nodes, including every intermediate node and link, conforms to the user specified security requirements. In essence, the proposed approach views security as a distributed network resource and allocates it to each user efficiently, based on demand and dictated by the need.

A Mathematical Framework for Synthesizing Distributed Algorithms: For many military and civilian large-scale, real-world problems of interest, data is first acquired asynchronously, i.e. at irregular intervals of time, at geographically-dispersed sites, processed utilizing decision-making algorithms, and the processed data then disseminated to other appropriate sites. Examples include the battlefield where intelligence data is gathered from different sources and processed to yield decisions and the proposed intelligent vehicle highway system where incident-related data is acquired along the highways to determine the best routes for traffic. The traditional approach to such problems consists of designing a central entity which collects all data and executes a decision making algorithm sequentially to yield the decisions. Centralized decision making algorithms are slow and highly vulnerable to natural and artificial catastrophes. Recent literature contains proposals for asynchronous, distributed, decision making algorithms wherein local decision making at every site replaces the centralized decision making to achieve faster response, higher reliability, and greater accuracy of the decisions. Two key issues with asynchronous, distributed, decision making algorithms include the lack of a suitable mechanism to synthesize them for any given problem and a comparative analysis of the quality of their

decisions.

This paper proposes a mathematical framework, MFAD, based on the Kohn-Nerode distributed hybrid control paradigm, to describe a centralized decision-making algorithm and to synthesize from it a distributed decision-making algorithm. Ideally, the centralized control gathers all data from all entities in the system and utilizes them to arrive at the decisions and, as a result, the decisions are expected to be "globally" optimal. In truth, however, as the frequency of the sensor data increases and the environment gets larger, dynamic, and more complex, the decisions come into question. In the distributed decision-making system, the centralized decision-making is replaced by those of the constituent entities that aim at minimizing a Lagrangian, i.e. a local, non-negative cost criterion, subject to the constraints imposed by the global goal. Thus, computations are carried out locally, utilizing locally obtained data and appropriate information that is propagated from other sites. It is hypothesized that with each entity engaged in optimizing its individual behavior, asynchronously, concurrently, and independent of other entities, the distributed system will achieve "global" optimal behavior. The distributed system is also expected to be more efficient, robust, and exhibit scalability. While it does not claim that decentralized systems may be synthesized for all centralized real-world systems, this paper implements both the centralized and distributed paradigms for a representative military battlefield command, control, and communication (C^3) problem. It also simulates them on a testbed of network of workstations for a comparative performance evaluation of the centralized and decentralized paradigms in the MFAD framework. Performance results indicate that the decentralized approach consistently outperforms the centralized approach. The paper aims at developing a quantitative evaluation of the quality of decisions under the decentralized paradigm and, to achieve this goal, it introduces a fundamental concept, embodied through a hypothetical entity termed "Perfect Global Optimization Device (PGOD)," that generates perfect or ideal decisions. PGOD possesses perfect knowledge, i.e. the state information of every entity of the entire system unaffected by time and delay, at all times. PGOD utilizes the same decision-making algorithm as the centralized paradigm and generates perfect globally-optimal decisions which, though unattainable, provide a fundamental and absolute basis for comparing the quality of decisions. Simulation results reveal that the quality of decisions in the decentralized paradigm are superior to those of the centralized approach and that they closely track PGOD's decisions.

LIST OF PUBLICATIONS

- Bienenstock E. and Geman S. 1995. Compositionality in Neural Systems. in The Handbook of Brain Theory and Neural Networks M.A. Arbib Ed. Bradford Books/MIT Press pp. 223–226.

- Bienenstock E. Geman S. and Potter D. 1997. Compositionality MDL Priors and Object Recognition. In: *Advances in Neural Information Processing Systems 9* M. Mozer M. Jordan and T. Petsche eds. Bradford/MIT pp 838–844.

- Chi Z. and S. Geman. (in press) Estimation of probabilistic context-free grammars. *Computational Linguistics.*

- Dupuis, P., Grenander, U., and Miller, M. *Variational Problems on Flows of Deffeomorphisms for Image Matching,* to appear in Quarterly of Applied mathematics, 1998.

- Geman S. D. Potter and Z. Chi. 1998. Composition systems. Technical report Division of Applied Mathematics Brown University.

- Geman S. 1998. Invariant binding in composition systems. Technical report Division of Applied Mathematics Brown University.

- Gidas, B. *Metropolis-type Monte Carlo Simulation Algorithms and Simulated Annealing,* in Topics in Contmporary Probability and Its Applications, CRC Press 1995, editor: J. L. Snell.

- Gidas, B., and Zelic, A. *Hierarchical Syntactic Models for Object Recogntion* Proc. of the 13-th International Conference on Digital Signal Processing, Santorini, Greece, 1997, pp. 1117-1121.

- Gidas, B., Robertson, C., and Almeida, M. *Object Tracking in Cluttered Enviroments via the Monte Carlo Filter,* in preparation.

- Joshi, S.C., Wang, J., Miller, M., Van Essen, D., and Grenander, U. *On the Differential Gemetry of the Cortical Surface,* In Proc. of SPIE's 1995 geometric Methods in Applied Imaging, San Diego, CA,, July 9-14, 1995.

- Miller, M., Khaneja, N., and Grenander, U. *Dynamic Programming Generation of Gedesics and Salci on Brain Surfaces,* to appear in IEEE PAMI 1999.

- Mumford, D., Zhu, S., and Gidas, B. *Stochastic Models for Generic Images,* Preprint, Brown University, 1998.

- Raphael, C., and Geman, S. *Mine Detection and Classification in Shallow Water Imagery.* Technical Report, Division of Applied Mathematics, Brown University, 1995.

- Robertson, C. *Tracking of Objects from Image Sequences Using Lagrangian Dynamics and Nonlinear Filtring,* Ph.D. Thesis, Division of Applied Mathematics, Brown University, Providence, R.I., 1998.

- Zelic, A. *A Stochastic Hierarchical Model and Nonparametric Statistical Approach of Object Recognition,* Ph.D. Thesis, Division of Applied Mathematics, Brown University, Providence, R.I., 1996

- H.J.(Jerry) Schumacher and Sumit Ghosh, "A Fundamental Framework for Network Security," Journal of Networks and Computer Applications, Vol. 20, No. 3, July 1997, pp. 305-322.

- Kwun Han and Sumit Ghosh, "A Comparative Analysis of Virtual Versus Physical Process-Migration Strategies for Distributed Modeling and Simulation of Mobile Computing Networks," ACM Wireless Network Journal – The Journal of Mobile Communication, Computation, and Information, Vol. 4, No. 5, 1998, pp. 365-378, Baltzer Science Publishers, Red Bank, NJ.

- Sumit Ghosh, Qutaiba Razouqi, H.J.(Jerry) Schumacher, and Aivars Celmins "A Survey of Recent Advances in Fuzzy Logic in Telecommunications Networks and New Challenges," IEEE Transactions on Fuzzy Systems, Vol. 6, No. 3, August 1998, pp. 443-447.

- Ricardo Citro, Sumit Ghosh, and Peter G. Churgin, "A Fundamental Metric for Continuity of Care: Modeling and Performance Evaluation," IEEE Transactions on Information Technology in Biomedicine, Vol. 1, No. 3, September 1997, pp. 189-204.

- Tony Lee and Sumit Ghosh, "Stability of RYNSORD, a Decentralized Algorithm for Railway Networks, Under Perturbations," Revised Version Submitted for possible publication in the IEEE Transactions on Vehicular Technology, January 1998.

- Qutaiba Razouqi, Tony Lee, and Sumit Ghosh, "A Guaranteed-No-Cells-Dropped Buffer Management Scheme with Selective Blocking for Cell-Switching Networks," Computer Communications, Vol 21. No. 10, July 15, 1998, pp. 930-946, Elsevier Science Publishers, B.V.

- H.J.(Jerry) Schumacher and Sumit Ghosh, "A Fundamental Framework for Network Security Towards Enabling Security on Demand in ATM Network," Computers & Security, Official Journal of the IFIP Technical Committee 11, September 1998. Elsevier Advanced Technology Publishers.

- H.J.(Jerry) Schumacher, Tony Lee, and Sumit Ghosh, "Top Security Traffic and the Public ATM Network Infrastructure," Journal of Information Systems Security, Vol. 7, No. 4, Winter 1999 (December), pp., Auerbach Publications.

- Tony Lee, S. Ghosh, J. Lu, X. Ge, A. Nerode, and W. Kohn, "A Mathematical Framework for Asynchronous, Distributed, Decision-Making Systems with Semi-Autonomous Entities: Algorithm Synthesis, Simulation, and Evaluation," Revised version submitted for possible publication in IEEE Transactions on Systems, Man, and Cybernetics Part B: Cybernetics, April 1998.

- Sumit Ghosh, Tony Lee, H.J. (Jerry) Schumacher, and Seong-Soon Joo, "A Frame of Reference for the Performance Evaluation of Asynchronous, Distributed Decision-Making Algorithms," Submitted for possible publication in the IEEE Transactions on Computers, May 1998.

- Qutaiba Razouqi and Sumit Ghosh, "Performance Analysis of Fuzzy Thresholding-Based Buffer Management for a Large-Scale Cell-Switching Network," Submitted to IEEE Transactions on Communications, February 1998.

- Seong-Soon Joo and Sumit Ghosh, "Generalized End-to-End Performance Estimation of ATM Networks Using Combined Simulation Model and Neural Network," Proceedings of the IEEE Globecom'98 Conference, Session 116.7, Genetic Algorithms, pp. 116.7.1 - 116.7.5, Sydney, Australia, November 8-12, 1998.

- Sumit Ghosh, "New Principles for the Design of Programming Languages and Software Execution Environments for the 21st Century," Proceedings of the Advanced Software Technology in the 21st Century and Related Topics in Computer Applications – The 1998 Stephen S. Yau International Workshop, Westin Hotel, Santa Clara, CA., March 8-9, 1998.

- Seong-Soon Joo and Sumit Ghosh, "A New Metric Towards Comprehensive Performance Evaluation of ATM Networks," Proceedings of the IEEE ATM '98 Conference, George Mason University, VA, May 26-29, 1998.

LIST OF PAERTICIPANTS

- Basilis Gidas, Professor
- Sumit Ghosh, Professor
- Stuart Geman, Professor
- Ulf Grenander, Professor
- Donald McClure, Professor
- Elie Bienenstock, Senior Visitor
- Antonio Zelic, Ph.D. 1996
- Kevin Kochanek, Ph.D., 1997
- Larisa Matejic, Ph.D. 1997
- Chris Robertson, Ph.D. 1998
- Zhiyi Chi, Ph.D. 1998
- Shih-Hsien Huang, Ph.D. candidate
- Daniel Potter, Ph.D. candidate